# Escaping the Simulation: Systematic Concept Expansion via Semantic Operators

Fred Zimmerman
Nimble Books LLC
wfz@nimblebooks.com

December 2025

**Abstract**

Consider every corpus as a simulation—a compressed, selective model of a richer underlying reality. When analysts write strategic assessments or scientists publish papers, they're constructing simplified models that capture certain features, compress others, and omit what seems irrelevant or simply wasn't thought of. We present a framework for *escaping the simulation*: using patterns within a corpus to infer concepts that exist in the fuller reality but were compressed away or never articulated.

The framework follows a principled pipeline: Corpus $\rightarrow$ Extract($type$) $\rightarrow$ Augment($knowledge$) $\rightarrow$ Define($operators$) $\rightarrow$ Apply($ops, concepts$) $\rightarrow$ Filter($threshold$). We demonstrate it through construction of a strategic lexicon from 821 national security documents, generating 16 novel doctrines that occupy genuine gaps in strategic vocabulary. The key insight: simulations leak information about what they're simulating. The concepts a corpus *does* contain imply other concepts it *doesn't*.

**Keywords:** concept expansion, semantic embeddings, RKHS, strategic doctrine, escaping simulations

## 1 Introduction

Consider every corpus as a simulation. When analysts, scientists, or strategists write documents, they're constructing a simplified model of a vastly more complex reality. The simulation captures certain features deemed important, compresses others, and omits entirely what seems irrelevant or what the authors simply didn't think to include.

This simulation is useful—that's why we build it. But it's not reality. The map is not the territory.

When we analyze a corpus, we're studying the simulation. We find the concepts the simulators chose to articulate. We trace the relationships they made explicit. We learn the vocabulary they developed. But we remain inside the simulation, seeing only what it was designed to show us.

The interesting question: what exists in the underlying reality that the simulation doesn't capture? What concepts *should* be there but aren't? What ideas lie in the gap between the model and the world it models?

This paper presents a framework for escaping the simulation—for using the structure and patterns within a corpus to infer concepts that exist in the fuller reality but were compressed away or never articulated.

### 1.1 The Core Pipeline

The framework follows six stages:

$$\boxed{\text{Corpus} \xrightarrow{\text{Extract}} \text{Concepts} \xrightarrow{\text{Augment}} \text{Enriched} \xrightarrow{\text{Define}} \text{Operators} \xrightarrow{\text{Apply}} \text{Candidates} \xrightarrow{\text{Filter}} \text{Novel}} \quad (1)$$

More precisely:

$$\begin{aligned}
&\text{Extract}(C,\tau) && \mapsto \{c_1,\ldots,c_n\} && \text{concepts of type } \tau \text{ from corpus } C && (2)\\
&\text{Augment}(S,K) && \mapsto S' && \text{enrich } S \text{ with external knowledge } K && (3)\\
&\text{Define}(\mathcal{D}) && \mapsto \{T_1,\ldots,T_m\} && \text{operators appropriate to domain } \mathcal{D} && (4)\\
&\text{Apply}(T_i,c_j) && \mapsto c' && \text{transform concept } c_j \text{ via operator } T_i && (5)\\
&\text{Filter}(c',\theta) && \mapsto \{0,1\} && \text{accept if novelty}(c') > \theta && (6)
\end{aligned}$$

Each stage is independent and admits multiple implementations. The framework specifies *what* each stage must accomplish, not *how*—enabling domain-specific instantiation while preserving theoretical coherence.

### 1.2 Contributions

1. A general framework for concept expansion applicable across domains
2. Formal definitions of novelty and coherence in semantic space
3. A library of ten morphological operators for strategic concepts
4. Empirical validation through construction of a 121-entry strategic lexicon

## 2 The Framework

### 2.1 Stage 1: Corpus

Input: A document collection $C = \{d_1, \ldots, d_N\}$ representing a domain's current knowledge state.

The corpus *is* the simulation—a compressed model of some underlying reality. The documents represent what the simulators (authors, analysts, strategists) chose to articulate. Our goal is to use this simulation to infer what lies beyond it.

**Requirements:**

- Coverage: $C$ should represent the domain comprehensively
- Authority: Documents should be authoritative sources
- Temporality: Consider whether $C$ represents a point-in-time or longitudinal view

### 2.2 Stage 2: Extract$(C, \tau)$

Input: Corpus $C$, concept type specification $\tau$

Output: Set of concepts $S = \{c_1, \ldots, c_n\}$ of type $\tau$

"Concept" is deliberately flexible. The type $\tau$ might specify:

- **Terms**: Named entities, technical vocabulary, doctrines
- **Claims**: Causal assertions, predictions, arguments
- **Patterns**: Design patterns, argument structures, narrative arcs
- **Relations**: Entity relationships, dependencies, hierarchies

**Implementation options:**

1. *Manual extraction*: Expert annotation of corpus
2. *NLP extraction*: Named entity recognition, term extraction
3. *LLM extraction*: Prompt-based concept identification
4. *Hybrid*: Automated extraction with expert validation

### 2.3 Stage 3: **Augment**$(S, K)$

Input: Extracted concepts $S$, external knowledge source $K$

Output: Enriched concept set $S' \supseteq S$

This is where we begin escaping the simulation. Augmentation brings in concepts from the underlying reality that the corpus-builders didn't include—whether from adjacent domains, expert knowledge, or *other simulations of the same reality*.

External knowledge $K$ might include:

- Adjacent domains (strategic $\rightarrow$ economic concepts)
- Historical sources (contemporary $\rightarrow$ classical foundations)
- Cross-cultural perspectives (Western $\rightarrow$ Eastern doctrines)
- Domain ontologies and taxonomies

The augmented set $S'$ represents a less compressed view of reality than the original corpus alone.

### 2.4 Stage 4: **Define**$(\mathcal{D})$

Input: Domain specification $\mathcal{D}$

Output: Operator library $\mathcal{T} = \{T_1, \ldots, T_m\}$

Operators encode domain knowledge about what transformations produce meaningful results. We distinguish:

**Universal operators** (applicable across domains):

$$\text{Compose}(a, b) \rightarrow \text{compound concept} \tag{7}$$
$$\text{Negate}(a) \rightarrow \text{conceptual opposite} \tag{8}$$
$$\text{Specialize}(a, d) \rightarrow \text{apply } a \text{ to domain } d \tag{9}$$
$$\text{Abstract}(a) \rightarrow \text{generalize } a \tag{10}$$

**Domain-specific operators** (require domain expertise to define):

$$\text{Temporalize}(a) \rightarrow \text{add temporal dynamics} \tag{11}$$
$$\text{Adversarialize}(a) \rightarrow \text{opponent's perspective} \tag{12}$$
$$\text{Escalate}(a, \lambda) \rightarrow \text{move along intensity spectrum} \tag{13}$$

The quality of generated concepts depends critically on operator design. Poorly chosen operators produce garbage.

### 2.5 Stage 5: Apply$(T, c)$

Input: Operator $T \in \mathcal{T}$, concept $c \in S'$

Output: Candidate concept $c' = T(c)$

Application generates the candidate space. For $n$ concepts and $m$ unary operators, we generate $O(nm)$ candidates. Binary operators (composition) generate $O(n^2)$. Multi-step application produces exponential growth.

**Scaling strategies:**

1. *Prioritize*: Apply operators to central/important concepts first
2. *Sample*: Random or importance-weighted sampling
3. *Prune early*: Filter for coherence before measuring novelty
4. *Iterate*: Multiple passes with human review narrowing focus

### 2.6 Stage 6: Filter$(c', \theta)$

Input: Candidate concept $c'$, novelty threshold $\theta$

Output: Accept/reject decision

The filter stage separates genuinely novel concepts from noise. Two criteria:

**Novelty**: Distance from existing concept space

$$\text{novelty}(c') = \min_{c \in S'} d(\phi(c'), \phi(c)) \tag{14}$$

where $\phi$ is the embedding function and $d$ is semantic distance.

**Coherence**: Does $c'$ represent a valid concept?

- Semantic coherence: Consistent meaning, definable
- Domain coherence: Recognized by experts as plausible
- Compositional coherence: Parts combine sensibly

The *novelty-coherence tradeoff*: Very novel candidates are often incoherent; very coherent candidates are often not novel. The sweet spot is concepts that are *surprisingly* coherent—distant from the simulation's explicit content but still grounded in the underlying reality it models.

# 3 Theoretical Foundations

## 3.1 Semantic Embeddings as Feature Maps

We ground the framework in Reproducing Kernel Hilbert Space (RKHS) theory. A kernel $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ defines an implicit feature map $\phi : \mathcal{X} \to \mathcal{H}$ such that:

$$K(c_1, c_2) = \langle \phi(c_1), \phi(c_2) \rangle_{\mathcal{H}} \tag{15}$$

Modern sentence embeddings approximate this structure. If $\text{embed}(c) \in \mathbb{R}^n$ is the embedding, then:

$$K(c_1, c_2) = \text{embed}(c_1)^\top \text{embed}(c_2) \tag{16}$$

defines a valid linear kernel.

## 3.2 Operators as Transformations

Given operator $T$ on concepts, we seek induced operator $\hat{T}$ on $\mathcal{H}$ such that:

$$\phi(T(c)) \approx \hat{T}(\phi(c)) \tag{17}$$

This isn't always exact (concept operations aren't always linear in embedding space), but the approximation guides implementation.

## 3.3 Novelty as Distance from the Simulation

The simulation (corpus) provides a sparse sampling of conceptual space. Novelty measures distance from sampled points—how far a candidate lies from what the simulation explicitly contains:

$$\text{novelty}(c_{\text{new}}) = \min_{c \in S} \|\phi(c_{\text{new}}) - \phi(c)\|_{\mathcal{H}} \tag{18}$$

Using the kernel trick:

$$\|\phi(c_1) - \phi(c_2)\|^2 = K(c_1, c_1) - 2K(c_1, c_2) + K(c_2, c_2) \tag{19}$$

This allows novelty computation without explicit high-dimensional representations. Operators act as transformations that move us toward unsampled regions—portions of the underlying reality that the simulation compressed away.

# 4 Application: The Strategic Lexicon

We demonstrate the framework through construction of a comprehensive strategic lexicon.

## 4.1 Corpus (Stage 1)

**Input:** 821 national security documents spanning 1987–2024, drawing on datasets from Gannon [1] and the Military Doctrines project [2]:

- U.S. National Security Strategies: 12 documents

- NATO Strategic Concepts: 4 documents
- Chinese Defense White Papers: 10 documents
- Russian Military Doctrines: 5 documents
- EU Security Strategies: 6 documents
- Allied National Strategies: 784 documents from 50+ countries

### 4.2 Extract (Stage 2)

**Type specification:** Strategic doctrines—named concepts representing positions on security, force employment, and international relations.

**Method:** Hybrid LLM-assisted extraction with expert validation.

**Output:** 84 documentary doctrines including:

- Cold War vocabulary: *containment, deterrence, mutually assured destruction*
- Post-Cold War terms: *transformation, revolution in military affairs*
- Contemporary framing: *great power competition, integrated deterrence, gray zone*

### 4.3 Augment (Stage 3)

**External knowledge:** The Western strategic canon—12 foundational texts from Thucydides through Liddell Hart.

**Output:** 21 classical concepts added:

- From Thucydides: Dynamics later named *security dilemma* (Herz, 1950)
- From Clausewitz: *center of gravity, friction, fog of war*
- From Sun Tzu: *acme of skill* (winning without fighting)
- From Mackinder: *heartland theory*

Total enriched set: $|S'| = 105$ concepts.

### 4.4 Define (Stage 4)

Ten morphological operators designed for strategic doctrine:

### 4.5 Apply (Stage 5)

Systematic application to enriched concept set. Example applications:

**From documentary concepts:**

$$\text{Blend}(\text{deterrence}, \text{gray zone}, \text{resilience}) \rightarrow \textit{Adaptive Deterrence}$$
$$\text{Domain}(\text{containment}, \text{space}) \rightarrow \textit{Spatial Containment}$$
$$\text{Culture}(\text{collective defense}, \text{US} \rightarrow \text{China}) \rightarrow \textit{Civilizational Defense}$$

| Operator | Formula | Effect |
|---|---|---|
| Invert | $-f$ | Counter-doctrine |
| Temporal | $(1-\alpha)f + \alpha c_{\text{era}}$ | Era translation |
| Escalate | $f + \lambda v_{\text{esc}}$ | Intensity shift |
| Culture | $f - c_s + c_t$ | Cultural translation |
| Domain | $(1-\alpha)f + \alpha c_{\text{dom}}$ | Domain projection |
| Actor | $(1-\alpha)f + \alpha c_{\text{actor}}$ | Actor transformation |
| Blend | $\sum_i w_i f_i$ | Synthesis |
| Constrain | $f - \sum_j \langle f, c_j \rangle c_j$ | Element removal |
| Intensify | $c + \alpha(f - c)$ | Commitment amplification |
| Negate | $c_{\text{all}} - f$ | Explicit rejection |

Table 1: RKHS morphological operators for strategic doctrine

**From classical concepts:**

$$\text{Blend}(\text{friction, cyber warfare}) \to \textit{Digital Friction}$$
$$\text{Blend}(\text{indirect approach, gray zone}) \to \textit{Hybrid Indirect Approach}$$
$$\text{Domain}(\text{heartland theory, space}) \to \textit{Space Heartland}$$

### 4.6 Filter (Stage 6)

**Novelty threshold:** Kernel distance $> 0.3$ from nearest existing concept.

**Coherence check:** LLM-based evaluation ("Is this a coherent strategic concept?") plus manual expert review.

**Output:** 16 novel doctrines accepted:

- 8 from documentary sources
- 8 from classical sources

### 4.7 Summary

| Stage | Count | Description |
|---|---|---|
| Corpus | 821 | Documents analyzed |
| Extract | 84 | Documentary doctrines |
| Augment | +21 | Classical concepts added |
| Define | 10 | Operators designed |
| Apply | ∼500 | Candidates generated |
| Filter | 16 | Novel doctrines accepted |
| **Final** | 121 | Total lexicon entries |

Table 2: Pipeline statistics for strategic lexicon construction

## 5 Generalization Across Domains

The framework applies beyond strategic doctrine. We sketch applications:

### 5.1 Scientific Hypothesis Generation

- **Corpus:** Published papers in research area
- **Extract type:** Causal claims, mechanisms, phenomena
- **Augment:** Adjacent fields, historical papers
- **Operators:** Compose, Condition, Negate
- **Output:** Testable hypotheses not yet explored

### 5.2 Legal Argument Development

- **Corpus:** Case law, statutes, legal scholarship
- **Extract type:** Legal principles, argument structures
- **Augment:** Other jurisdictions, legal philosophy
- **Operators:** Distinguish, Extend, Analogize
- **Output:** Novel but defensible legal arguments

### 5.3 Product Innovation

- **Corpus:** Patents, product descriptions, user research
- **Extract type:** Features, user needs, technical solutions
- **Augment:** Analogous products, emerging technologies
- **Operators:** Combine, Transfer, Invert
- **Output:** Systematic ideation beyond brainstorming

## 6 Limitations and Future Work

### 6.1 Current Limitations

**Operator design requires expertise:** The framework doesn't specify *which* operators to use—that requires domain knowledge.

**Coherence is hard to automate:** LLM-based coherence checking is imperfect. Some nonsense passes; some good ideas fail.

**Semantic distance isn't sufficient:** Two concepts can be distant in embedding space but not interestingly different.

**Combinatorial explosion:** Even modest concept sets generate huge candidate spaces.

### 6.2 Future Directions

**Learned operators:** Can we learn domain-appropriate operators from examples of conceptual innovation?

**Better coherence models:** Train models specifically for conceptual coherence evaluation.

**Interactive exploration:** Tools for experts to navigate candidate space with feedback loops.

**Cross-domain transfer:** Do operators transfer across domains? Can we build a taxonomy?

## 7 Conclusion

Consider every corpus as a simulation—a compressed, selective model of a richer underlying reality. We presented a framework for escaping that simulation: Corpus → Extract → Augment → Define → Apply → Filter.

The key insight is that simulations leak information about what they're simulating. The concepts a corpus *does* contain imply other concepts it *doesn't*. The relationships it makes explicit suggest relationships it leaves tacit. The vocabulary it develops points toward vocabulary gaps. By making these implications systematic—through principled operators and novelty measurement—we can reconstruct portions of the underlying reality that the simulation doesn't directly represent.

The framework is most powerful when combined with domain expertise—expertise in understanding what the simulation was designed to capture, what it likely compresses or omits, and what operators might expose those hidden structures. It doesn't replace human creativity but augments it, systematically exploring regions of conceptual reality that the simulation renders invisible.

As language models and embedding techniques improve, our ability to escape simulations grows. Better embeddings mean more accurate maps of conceptual space. Better generation models mean more sophisticated reconstruction of what's missing. Better evaluation models mean more reliable detection of when we've found something real versus when we've wandered into noise.

The goal isn't to automate creativity—it's to develop systematic methods for inferring what lies beyond the boundary of any particular simulation. Every corpus shows us part of the world. The Concept Expansion Framework helps us see the rest.

## Acknowledgments

This framework was developed during construction of a strategic lexicon for national security analysis. The strategic lexicon is published separately as *The Strategic Lexicon: A Dictionary of National Security Concepts* (Victory Conditions Ventures, December 2025).

## References

[1] Gannon, J.A. (2024). A Text-as-Data Approach to Measuring Interstate Relations Using National Security Strategies. Working paper.

[2] Military Doctrines Dataset. (2025). *Military Doctrines*. Retrieved December 11, 2025, from https://militarydoctrines.com/

[3] Herz, J.H. (1950). Idealist Internationalism and the Security Dilemma. *World Politics*, 2(2), 157–180.

[4] Allison, G. (2017). *Destined for War: Can America and China Escape Thucydides's Trap?* Houghton Mifflin Harcourt.

[5] Kennedy, P. (1987). *The Rise and Fall of the Great Powers*. Random House.

[6] von Clausewitz, C. (1832/1976). *On War*. Princeton University Press.

[7] Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *EMNLP 2019*.